

POWER CONSERVATION IN A SERVER CLUSTER

BACKGROUND

5 1. Field of the Present Invention

The present invention relates generally to the field of data processing networks and more particularly to a server cluster and method in which cached files on powered-down servers remain accessible to the active servers to conserve power consumption without substantially impacting response performance.

10 2. History of Related Art

In the field of network computing, network services are frequently provided using a server cluster in which multiple server services are connected together in a local area network configuration. Referring to FIG 2, a conceptualized illustration of an example in which Web based services are provided by a server cluster 118 is presented. Server cluster 118 includes a set of servers 120-123 connected together in a local area network. Server cluster 118 may be configured such that a single universal resource indicator is serviced by any of the servers in the cluster. In other words, clients use a single network address when directing requests to the cluster, and any client request can be serviced by any of the servers in the cluster. In this illustration, a series of client requests, each indicated by a letter in reference numeral 119, are received by server cluster 118. Each letter in the series corresponds to the target file of the respective request. Each letter might, for example, identify a unique universal resource locator (URL). The series of requests 119 may include multiple requests for the same target file. The 20 target file corresponding to letter "A," for example, has been requested twice in the illustrated series.

25 The server cluster 118 includes a request distributor 126 that implements locality aware request distribution (LARD). In a LARD configuration, client requests are routed to respective servers based at least in part on the requested target file. Routing requests in this manner

attempts to distribute loading over the servers and to maximize efficiency by sending requests to the server that most recently serviced a request for a specific target file. In the illustrated example, requests for target files A and E are routed to server 120, requests for target files B and F are routed to server 121, requests for target files C and G have been routed to server 122, and 5 requests for target files D and H have been routed to server 123.

Each server is shown as including a disk cache or file cache **124**. File cache **124** typically comprises a portion of the respective server's volatile system memory and includes files or other data objects recently retrieved from disk storage **125**. Thus, the file cache **124** of server **120** is shown as including requested files A and E while file cache **124** of server **121** includes requested files B and F and so forth. Data stored in file caches **124**, of course, is retrieved faster than data stored in disk storage **125**.

While distributing requests across multiple servers in a server cluster may provide performance advantages, it may also increase operating costs. The power consumed by each active server in a server cluster has motivated cluster designs in which servers are dynamically powered-down in response to the varying loads of client requests. At times when request loading is low, fewer servers are required to handle the client requests with adequate response performance while, at peaks loads, all servers may be required to deliver acceptable performance.

The advantages offered by LARD may, unfortunately, be inconsistent with the advantages offered by dynamic power loading. More specifically, by dynamically powering down servers to account for reduced loading, a server cluster may negatively impact performance by powering down one or more servers that contain cached copies of recently requested target files. Referring to FIG 3, an illustration of the performance penalty that occurs when one or more servers is powered-down in a LARD configured server cluster is presented. In this illustration, servers **120** and **121** have been powered-down by the server cluster in an effort to conserve power during a time of reduced traffic. Prior to power-down, server **120** had been servicing requests for target files A and E (as in FIG 2) and server **121** had been servicing requests for target files B and F. Accordingly, the file caches **124** of the remaining active servers **122** and **123** do not contain cached copies of target files A, E, B, and F. When a subsequent request for any of these target files is received, the server cluster must direct the request to one of

the active servers. Because the file caches of the active servers do not include copies of target files A, B, E, and F, subsequent requests for these files have to be serviced from disk storage 125. It would be desirable to implement a server cluster that benefited from dynamic activation of servers and locality aware request distribution without suffering the disk access performance 5 penalty described above.

SUMMARY OF THE INVENTION

10 The problems identified above are addressed by a data processing system and method in which servers on a server cluster are selectively deactivated based upon loading or some other parameter. A server is deactivated by completing any pending client requests and then transitioning the server's processor(s) to a low power state. The deactivated server continues to provide power to its system memory and network interface card. Prior to deactivation, a server enables the other servers to access its file cache, such as by broadcasting a directory of its disk 15 cache or file cache to other servers on the network. The network interface cards are configured with direct memory access capability. While a server is deactivated, the remaining active servers on the network may access files or other data objects within the deactivated server's file cache over the network through the deactivated server's network interface card.

20

BRIEF DESCRIPTION OF THE DRAWINGS

Other objects and advantages of the invention will become apparent upon reading the following detailed description and upon reference to the accompanying drawings in which:

25 FIG 1 is a block diagram of selected elements of a data processing network according to one embodiment of the present invention;

FIG 2 is a conceptualized illustration of a server cluster employing locality aware request distribution;

FIG 3 is a conceptualized depiction of the server cluster of FIG 2 illustrating operation of

the cluster after one or more servers have been powered-down to conserve energy;

FIG 4 is a block diagram of selected features of a server according to one embodiment of the invention;

FIG 5 is a block diagram of selected features of a network interface card according to one embodiment of the invention; and

FIG 6 is a conceptualized depiction illustrating operation of a server cluster according to the present invention.

While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description presented herein are not intended to limit the invention to the particular embodiment disclosed, but on the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
259
260
261
262
263
264
265
266
267
268
269
269
270
271
272
273
274
275
276
277
278
279
279
280
281
282
283
284
285
286
287
288
289
289
290
291
292
293
294
295
296
297
298
299
299
300
301
302
303
304
305
306
307
308
309
309
310
311
312
313
314
315
316
317
318
319
319
320
321
322
323
324
325
326
327
328
329
329
330
331
332
333
334
335
336
337
338
339
339
340
341
342
343
344
345
346
347
348
349
349
350
351
352
353
354
355
356
357
358
359
359
360
361
362
363
364
365
366
367
368
369
369
370
371
372
373
374
375
376
377
378
379
379
380
381
382
383
384
385
386
387
388
389
389
390
391
392
393
394
395
396
397
398
399
399
400
401
402
403
404
405
406
407
408
409
409
410
411
412
413
414
415
416
417
418
419
419
420
421
422
423
424
425
426
427
428
429
429
430
431
432
433
434
435
436
437
438
439
439
440
441
442
443
444
445
446
447
448
449
449
450
451
452
453
454
455
456
457
458
459
459
460
461
462
463
464
465
466
467
468
469
469
470
471
472
473
474
475
476
477
478
479
479
480
481
482
483
484
485
486
487
488
489
489
490
491
492
493
494
495
496
497
498
499
499
500
501
502
503
504
505
506
507
508
509
509
510
511
512
513
514
515
516
517
518
519
519
520
521
522
523
524
525
526
527
528
529
529
530
531
532
533
534
535
536
537
538
539
539
540
541
542
543
544
545
546
547
548
549
549
550
551
552
553
554
555
556
557
558
559
559
560
561
562
563
564
565
566
567
568
569
569
570
571
572
573
574
575
576
577
578
579
579
580
581
582
583
584
585
586
587
588
589
589
590
591
592
593
594
595
596
597
598
599
599
600
601
602
603
604
605
606
607
608
609
609
610
611
612
613
614
615
616
617
618
619
619
620
621
622
623
624
625
626
627
628
629
629
630
631
632
633
634
635
636
637
638
639
639
640
641
642
643
644
645
646
647
648
649
649
650
651
652
653
654
655
656
657
658
659
659
660
661
662
663
664
665
666
667
668
669
669
670
671
672
673
674
675
676
677
678
679
679
680
681
682
683
684
685
686
687
688
689
689
690
691
692
693
694
695
696
697
697
698
699
699
700
701
702
703
704
705
706
707
708
709
709
710
711
712
713
714
715
716
717
718
719
719
720
721
722
723
724
725
726
727
728
729
729
730
731
732
733
734
735
736
737
738
739
739
740
741
742
743
744
745
746
747
748
749
749
750
751
752
753
754
755
756
757
758
759
759
760
761
762
763
764
765
766
767
768
769
769
770
771
772
773
774
775
776
777
778
779
779
780
781
782
783
784
785
786
787
788
789
789
790
791
792
793
794
795
796
797
797
798
799
799
800
801
802
803
804
805
806
807
808
809
809
810
811
812
813
814
815
816
817
818
819
819
820
821
822
823
824
825
826
827
828
829
829
830
831
832
833
834
835
836
837
838
839
839
840
841
842
843
844
845
846
847
848
849
849
850
851
852
853
854
855
856
857
858
859
859
860
861
862
863
864
865
866
867
868
869
869
870
871
872
873
874
875
876
877
878
879
879
880
881
882
883
884
885
886
887
888
889
889
890
891
892
893
894
895
896
897
897
898
899
899
900
901
902
903
904
905
906
907
908
909
909
910
911
912
913
914
915
916
917
918
919
919
920
921
922
923
924
925
926
927
928
929
929
930
931
932
933
934
935
936
937
938
939
939
940
941
942
943
944
945
946
947
948
949
949
950
951
952
953
954
955
956
957
958
959
959
960
961
962
963
964
965
966
967
968
969
969
970
971
972
973
974
975
976
977
978
979
979
980
981
982
983
984
985
986
987
988
988
989
989
990
991
992
993
994
995
996
997
997
998
999
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1088
1089
1089
1090
1091
1092
1093
1094
1095
1096
1097
1097
1098
1099
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1188
1189
1189
1190
1191
1192
1193
1194
1195
1196
1197
1197
1198
1199
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1288
1289
1289
1290
1291
1292
1293
1294
1295
1296
1297
1297
1298
1299
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1388
1389
1389
1390
1391
1392
1393
1394
1395
1396
1397
1397
1398
1399
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1488
1489
1489
1490
1491
1492
1493
1494
1495
1496
1497
1497
1498
1499
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1588
1589
1589
1590
1591
1592
1593
1594
1595
1596
1597
1597
1598
1599
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1688
1689
1689
1690
1691
1692
1693
1694
1695
1696
1697
1697
1698
1699
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1788
1789
1789
1790
1791
1792
1793
1794
1795
1796
1797
1797
1798
1799
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1839
1840

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT
OF THE PRESENT INVENTION

5 Generally speaking the invention contemplates a server cluster and method of operation in which the servers on the cluster may be selectively powered-down or deactivated to reduce energy consumption. When a server is powered up or active, direct memory accesses to the server's memory from other servers on the cluster are prevented. After a server is deactivated, the complete contents of the deactivated server's file cache are available to other servers on the cluster. Before a server is deactivated, it may broadcast a directory of its file cache to the other servers in the cluster. When a server deactivates, it transitions its processor or processors to a low power state while maintaining power to its network interface card and system memory. After the server is deactivated, client requests to files or other data objects in the deactivated server's file cache may be processed by an active server. The active server assigned to handle the request may retrieve the file from deactivated server's memory over the network thereby avoiding a time consuming retrieval of data from disk. In this manner, a server cluster can implement a selective power down scheme without substantially degrading performance.

10 Turning now to the drawings, FIG 1 depicts selected features of a data processing network 100 according to one embodiment of the invention. In the depicted embodiment, network 100 includes a client 102 connected to a wide area network 104. Client 102 typically includes a client application program such as a conventional web browser that is executing on a client device. The client device may comprise a desktop or laptop personal computer, a network computer or workstation, or another network aware device such as a personal digital assistant (PDA) or an Internet enabled phone. Although client 102 is illustrated as remotely connected to 15 server network 101 through the intervening WAN 104, other clients (not depicted in FIG 1) may comprise a device on the server network itself.

20 Wide area network 104 typically includes various network devices such as gateways, routers, hub, and one or more local area networks (LANs) that are interconnected with various media possibly including copper wire, coaxial cables, fiber optic cables, and wireless media.

Wide area network 104 may represent or include portions of the Internet.

In the depicted embodiment, a server network or server cluster 101 is connected to client 102 through a gateway 106 connected to wide area network 104. Server cluster 101 is typically implemented as a LAN that includes one or more servers 110 (four of which are shown). The servers 110 may be networked together over a shared medium such as in a typical Ethernet or Token ring configuration. The depicted embodiment of server cluster 101 is implemented as a switched or point-to-point network in which each server 110 is connected to a switch 108 that is capable of providing a dedicated connection between any pair of servers 110.

The servers 110 of server cluster require access to a persistent (non-volatile) storage medium such as a magnetic hard disk. Any server 110 may include its own internal disk and disk drive facilities. In an increasingly prevalent configuration, persistent storage is provided as a networked device or set of devices. Networked storage is identified in FIG 1 by reference numeral 114 and may be implemented as one or more network attached storage (NAS) devices, a storage area network (SAN) or a combination thereof.

Referring now to FIG 4, a block diagram of selected elements of server 110 are depicted. Server 110 includes one or more processors 140 that can access a system memory 142 via a system bus 141. Processors 140 are commercially distributed general purpose microprocessors such as the PowerPC® family of processors from IBM Corporation or an x86-type processor such as the Pentium® family of processors from Intel Corporation. A bus bridge 144 connects system bus 141 to a peripheral bus 146. The peripheral bus is typically implemented as a Peripheral Components Interface (PCI) compliant bus to which one or more PCI compliant peripheral devices may be connected. In the depicted embodiment, a network interface card (NIC) 148 connected to peripheral bus 146 provides a connection between server 110 and an external network such as the local area network of server cluster 101. NIC 148 is configured with direct memory access capability that enables it to access the system memory 142 of server 110 without invoking or requiring the assistance of processor 140. In addition, server 110 is enabled to power down the processor(s) 140 without powering off NIC 148 or system memory 142.

Referring to FIG 5, selected elements of the NIC 148 according to one embodiment of the invention are depicted. In the depicted embodiment, NIC 148 includes an embedded controller

150 that can access a local memory 152 and a read only memory (ROM) element 153. NIC 148 further includes a DMA unit 156 connected to controller 150. Portions of the invention and NIC 148 may be implemented as a set of computer executable instructions (software) stored on a computer readable medium such as ROM 153. ROM 153 may comprise a flash memory card or 5 an electrically erasable programmable ROM (EEPROM).

Controller 150 is connected to a network interface unit 158 that provides buffering and buffering logic between the server cluster network and controller. Similarly, a peripheral bus interface unit 154 provides appropriate buffering between NIC controller 150 and peripheral bus 146. DMA unit 156 includes a set of DMA registers that are under the programmable control of controller 150. These registers may include a DMA command register for controlling DMA configuration parameters such as the active level of DMA request and acknowledge signals, a DMA mode register defining the DMA transfer mode, address increment/decrement setting, and other DMA mode attributes, a status register indicating whether a DMA request is pending, and one or more DMA address registers. When NIC 148 is instructed to initiate a DMA cycle, controller 150 sets the registers in DMA unit including the address register(s). Controller 150 may then instruct DMA unit 156 to execute the DMA cycle and DMA unit 156 responds by retrieving data from system memory 142 at the address defined by the address registers.

Server cluster 101 is implemented with a power management capability that monitors the server cluster loading and dynamically adjusts the number of servers 110 that are active at any 20 time based on the loading. When the server cluster traffic is low, fewer servers are required to be active. Thus, server cluster 101 may de-activate one or more servers 110 in response to low levels of network traffic. In one embodiment of the invention, server cluster 101 de-activates a server 110 by transitioning the server's processor(s) to a low power state. Commercially distributed microprocessors typically include software commands or resets that transition the 25 processor into a reduced power consumption state. Such states may be referred to as sleep mode, standby mode, and a variety of other terms. In any of these low power states, the processors 140 draw substantially less current than in the normal operating state.

When server cluster 101 determines that cluster traffic justifies the deactivation of one or more servers 110, the server(s) 110 selected for de-activation are transitioned to a power state in

which the processors 140 are in a low power state while power remains supplied to NIC 148, bus bridge 144, peripheral bus 146, system bus 141, and system memory 144. Server cluster 101 may send a message to a server 110 informing the server that has been selected for deactivation.

When a server 110 is selected for de-activation, the server will first complete any active requests that are pending. Upon completing any pending requests, a server selected for deactivation places its processors 140 into low power state. The selected server remains in this state until it is instructed to return to full power by the request distributor 126.

Each server 110 on server cluster 101 may implement a simple protocol that prevents other servers from accessing the contents of its memory when it is powered up, but allows the other servers to access the contents of its file cache when it is deactivated. In one embodiment, a selected server 110 may broadcast to each server 110 in server cluster 101 a directory of the server's file cache. Each server 110 that receives such a broadcast message will store the received directory in its memory. After a server 110 has broadcast its file cache directory to the server cluster, the server may enter a low power state by issuing an appropriate command or interrupt to the server's processor 140. If server 110 is subsequently instructed to return to full power, it first broadcasts to each server 110 in server cluster 101 a message revoking access to its file cache. When each server 110 has acknowledged receipt of this message, server 110 can begin accepting and processing new client requests.

After a server 110 has been deactivated, client requests are no longer routed to the server. If another server, however, receives a client request for a file or other data object that was in the file cache of a de-activated server, the server 110 receiving the request will determine from its copy of the de-activated server's file cache directory that the de-activated server's system memory has a copy of the requested file. The server 110 processing the client request will then send a message to the de-activated server to retrieve the requested file. The de-activated server's NIC 148 will then perform a DMA retrieval of the requested file and forward it back to the server handling the request.

Referring to FIG 6, operation of an embodiment of server cluster 101 is illustrated in a conceptualized format analogous to the illustration of the operation of the conventional server cluster 118 presented in FIG 3. Similar to FIG 3, FIG 6 depicts a series of clients requests that

are received by server cluster 101. Server cluster 101 includes a request distributor 116 analogous to request distributor 126 of server cluster 118 that forwards each of the received requests to a server in the cluster. As depicted in FIG 6, server 110 has been deactivated at a time when target files A and E are in its file cache 122 while server 111 has been deactivated at a time when target files B and F are in its file cache 122. Prior to deactivation, directories of these file caches have been broadcast to servers 112 and 113. When a subsequent request for a target file contained in the file caches 122 of deactivated servers 110 or server 111 are received by server cluster 101 and routed by request distributor 116 to one of the remaining active servers (122 or 123), the active server handling the request retrieves the requested target files from the system memory of the appropriate deactivated servers using the DMA capability of the deactivated server's NIC. In this manner, response performance is improved by reducing accesses to disk storage 114. Thus, FIG 6 illustrates a retrieval of target files A and E by server 112 from file cache 122 of deactivated server 110 and a retrieval of target files B and F by server 113 from file cache 122 of deactivated server 111. This DMA retrieval of files from the system memories of deactivated servers on the cluster is contrasted to the retrieval of requested files from disk storage as depicted in FIG 3 when a requested file resides in the file cache of a deactivated server.

It will be apparent to those skilled in the art having the benefit of this disclosure that the present invention contemplates a method of implementing a power managed server cluster without unnecessarily increasing the number of data accesses to disk storage. It is understood that the form of the invention shown and described in the detailed description and the drawings are to be taken merely as presently preferred examples. It is intended that the following claims be interpreted broadly to embrace all the variations of the preferred embodiments disclosed.